



NASA Center for Climate Simulation Success Story

Supercomputer Upgrades to Give Scientists Ability to Improve Forecasting

July 26, 2016

The Discover supercomputer saw a sizeable upgrade in May that will allow NASA Earth scientists to increase the resolution and capabilities of their models for more accurate predictions.

The NASA Center for Climate Simulation (NCCS) added a new scalable compute unit, SCU13, and 10 petabytes of storage to Discover. SCUs are supercomputers in their own right that are plugged into the supercomputer cluster through high-speed network connections. Like three of the four existing units, SCU13 is an SGI Rackable system.

“Because this was the exact same hardware from the exact same vendor, we had the opportunity with this upgrade to more tightly integrate it with some of the existing hardware, specifically SCU10,” said Bruce Pfaff, Discover’s lead system administrator.

SCU10 has 1,080 nodes, and SCU13 has 648 nodes. With each node housing 28 Intel E5-2697v3 (Haswell) processor cores, this integration creates a single system with 1,728 nodes and 48,384 cores.

“Some of our modeling users who run jobs that are 10,000, 20,000, or 30,000 cores can now run an even larger job than they could have before if we didn’t tightly integrate,” Pfaff said.



NCCS and SGI staff bring in and connect SCU13 to the Discover supercomputer. Photos by Jarrett Cohen, GSFC.

The integrated SCU10+SCU13 unit is capable of nearly 2 petaflops peak computing, or 2,000 trillion floating-point operations per second. It represents more than half of Discover.

To put the new computing power into context, assume the Earth’s population to be 7.4 billion people. In one second this system can calculate as much as every single person on Earth multiplying two numbers together every second for 75.5 hours straight.

Discover is divided into sections of about 600 nodes each, based on high-speed network switches. For users who need more than 600 nodes to run their models, it has been challenging to schedule the resources. Integrating SCU13 into SCU10 provides an area with over 1,700 nodes available to run a single application for users who need it. The integration opens up the potential to run a single application to more than three times what the NCCS has enabled in the past.

“What we’re doing is continuing to push the envelope of the current climate research models within NASA,” said Daniel Duffy, NCCS high-performance computing lead. “This has given us the ability to sit down with some of our primary customers and start them down the path of ‘Okay, now we have these resources, what can you do with them?’”

The Discover supercomputer is usually upgraded every year with either new technology or more of the same technology. Since the SCU13 upgrade was tightly integrated with SCU10, this upgrade used the same technology.

“One good reason to use that same architecture is because your users already know how to use it,” said Duffy. “That consistent architecture, first off, is very cost effective and, secondly and more importantly, the users can take advantage of it immediately without having to change their application or workflow.”

NCCS users include scientists in the Global Modeling and Assimilation Office (GMAO), the Goddard Institute for Space Studies (GISS), the NASA-Unified Weather Research and Forecasting (NU-WRF), and other general research and analysis groups that receive NASA funding.



The integration of SCU13 with SCU10 creates a system with 1,728 nodes and 48,384 cores. Photo by Naema Ahmed, GSFC.

“We couldn’t do anything we do without Discover,” said William Putman, who leads the GMAO model development group.

Planning models and projects at the GMAO requires assessing the technology available to them. The computing power available determines the advancement of the model. “Our models constantly push for more computing and faster networks,” said Putman. “Every time we get a new SCU, it gives us an opportunity to expand, explore resolutions we haven’t been to before, and add more complex models that take more time to compute.”

The SCU10+SCU13 combination allows the GMAO to focus on testing model resolutions four to eight times higher than before, according to Putman. The GMAO is currently evaluating their analysis and forecast system at 12.5-kilometer resolution.

Higher resolution enables models to use more physical equations and less estimation, yielding more details for scientists.

“We plan to use SCU13 for near-real-time operation, preparing a high-resolution reanalysis product, and using those advancements for coupled and ensemble models,” said Putman. With coupled and ensemble models scientists can look at several factors at once and use the models for things such as seasonal forecasting.

The models can be used to make seasonal forecasts because they allow scientists to look at how the ocean and atmosphere change together. In this case, higher resolution would ultimately mean better weather and climate prediction.

“The NCCS keeps up with what we need, but we continue to push towards bigger and faster to expand our simulations,” said Putman.

Naema Ahmed, NASA Goddard Space Flight Center

Contacts

Dan Duffy
High-Performance Computing Lead
NASA Center for Climate Simulation
NASA Goddard Space Flight Center
daniel.q.duffy@nasa.gov
301.286.8830

Bruce Pfaff
Lead System Administrator, Discover Supercomputer
NASA Center for Climate Simulation
NASA Goddard Space Flight Center
bruce.e.pfaff@nasa.gov
301.286.8587

William Putman
Model Development Lead
Global Modeling and Assimilation Office
NASA Goddard Space Flight Center
william.m.putman@nasa.gov
301.286.2599

More Information

Discover Supercomputer
<http://www.nccs.nasa.gov/services/discover>

Global Modeling and Assimilation Office
<http://gmao.gsfc.nasa.gov>



NASA Center for Climate Simulation
High-Performance Science
<http://www.nccs.nasa.gov>